AD-A117 748 UNIVERSITY OF SOUTHERN CALIFORNIA LOS ANGELES DEPT O--ETC F/G 17/1
HIGHLY PARALLEL MODERN SIGNAL PROCESSING.(U)
FEB 82 S KUNG

UNCLASSIFIED NL

1 of 1

END
DATE
FILMED
8-82
DTIC

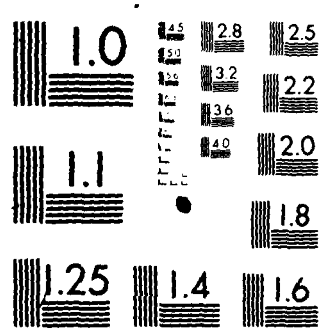

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

AD A117748

DTIC FILE COPY

6

ELECTRICAL ENGINEERING SYSTEMS DEPT.

UNIVERSITY OF SOUTHERN CALIFORNIA

Summary of Progress Report

HIGHLY PARALLEL MODERN SIGNAL PROCESSING

Sun-Yuan Kung
Principal Investigator
(213)743-6581

Covering Research Activity During the Period
1 March 1981 through 28 February 1982

Subcontracts:

Naval Ocean Systems Center
University of California, San Diego
Hughes Research Laboratory
Integrated Systems, Inc.
and
Stanford University

Sponsored by:

Office of Naval Research
Contract No. N00014 - 81 - K - 0191

DTIC
ELECTE
AUG 0 2 1982
S D
E

EE - Systems
University Park
Los Angeles, CA 90089 - 0272

This document has been approved for public release and sale; its distribution is unlimited.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: ONR - SRO REPORT NO. 1	2. GOVT ACCESSION NO.: AD-A117 748
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): HIGHLY PARALLEL MODERN SIGNAL PROCESSING	5. TYPE OF REPORT & PERIOD COVERED: ANNUAL PROGRESS REPORT 1 March 1981 to 28 Feb. 1982
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): Sun-Yuan Kung (Principal Investigator) and Subcontract Authors	8. CONTRACT OR GRANT NUMBER(s): N00014 - 81 - K - 0191
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Electrical Engineering - Systems, University of Southern California, Los Angeles, California 90089 - 0272	10. PROGRAM ELEMENT, ~~PROJECT~~, TASK AREA & WORK UNIT NUMBERS: NR SRO - 102/437
11. CONTROLLING OFFICE NAME AND ADDRESS: Office of Naval Research, 800 North Quincy Street, Arlington, VA 22217	12. REPORT DATE: 28 February 1982
	13. NUMBER OF PAGES: 48 pages
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): Office of Naval Research Western Regional Off., 1030 E. Green Street, Pasadena, CA 91106	15. SECURITY CLASS. (of this report): UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

APPROVED FOR RELEASE; DISTRIBUTION UNLIMITED.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

VLSI Signal processing theory and algorithms; VLSI Implementation
Parallel computing architectures/Toeplitz System solver
Wavefront Array Processor
Array processor testbed
Image Processing; Seismic Signal Processing

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes our research activities under SRO project, performed by University of Southern California, Naval Ocean System Center, Hughes Research Laboratories, Univ. of Calif., San Diego, Stanford University and Integrated Systems, Inc., for the period 1 March 1981 to 28 February 1982 under Contract No. N00014-81-K-0191 with the Office of Naval Research. This research activities have focussed on the VLSI signal processing theory and algorithms and the development of parallel computing architectures.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

S/N 0102-LF-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

A solution in today's VLSI research challenge lies in a cross-disciplinary research encompassing the areas of mathematics, algorithms, computers and applications. To this end, this report summarizes two parallel major research tasks: (1) Signal processing algorithm and theory - emphasizing spectral analysis and its applications; and (2) parallel computing structures - utilizing VLSI potential for high-speed signal processing.

Accession For	
NTIS GRA&I	X
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

DTIC COPY INSPECTED 2

SUMMARY OF PROGRESS REPORTS

ONR-SRO II PROJECT ON

HIGHLY PARALLEL MODERN SIGNAL PROCESSING

University of Southern California
Naval Ocean System Center
University of California, San Diego
Hughes Research Laboratory
Integrated Systems, Inc.
and
Stanford University

TABLE OF CONTENTS

ABSTRACT

This report describes the research activities performed by the University of Southern California, Naval Ocean System Center, Hughes Research Laboratories, University of California, San Diego, Stanford University and Integrated Systems, Inc., under the the SRO project, for the period 1 March 1981 to 28 February 1982 under contract No.: N00014-81-K-0191 with the Office of Naval Research. The research activities have focussed on the VLSI signal processing theory and algorithms and the development of parallel computing architectures.

A solution in today's VLSI research challenge lies in a cross-disciplinary research encompassing the areas of mathematics, algorithms, computers and applications. To this end, this report summarizes two parallel major research tasks: (1) Signal processing algorithm and theory - emphasizing spectral analysis and its applications; and (2) parallel computing structures - utilizing VLSI potential for high-speed signal processing.

This report describes our research activities under the SPO Project for the period March 1, 1981 to February 28, 1982. The research activities have focussed on the study of VLSI signal processing theory and algorithms and the development of parallel computing architectures.

With the rapidly growing microelectronics technology leading the way, modern signal processing is undergoing a major revolution. The availability of low cost, fast VLSI devices promises the practice of increasingly complex and sophisticated algorithms and systems. However, in conjunction with such promise, there is accompanied a new challenge of how to update the signal processing techniques so as to effectively utilize the large-scale computation capability. The answer to this challenge lies in a cross-disciplinary research encompassing the areas of mathematics, algorithms, computers and applications. To this end, two parallel major research tasks have been undertaken in the ONR-SPO research group;
(1) Signal Processing algorithm and theory - emphasizing spectral analysis and its applications; and (2) parallel computing structures - utilizing VLSI potential for high-speed signal processing.

In the area of signal processing theory and algorithms, significant work has been made on the following topics with special emphasis on high resolution spectral estimation: Recursive least squares ladder form algorithms (ISI, Stanford), adaptive least squares lattice (UCSD), Parallel Kalman filtering for both linear and non linear signal processing (ISI Stanford), adaptive notch filtering (USC), MEM, ARMA, and ML spectral estimation in 1-D and 2-D (USC, Stanford, and ISI), and Toeplitz approximation (USC). Parallel implementation of these algorithms has been a major consideration in their development.

In the complimentary area of parallel computing structures, both dedicated and flexible architectures have been developed for signal processing tasks and applications. Works in progress include: Toeplitz system solver using pipelined Levinson and implementation (USC and Hughes), programmable Wavefront Array Processor and data flow language for VLSI signal processing algorithms (USC), Systolic Arrays for real-time signal processing applications in Spectrum Analysis and Direction finding(NOSC), and implementation of Fadeev algorithm (Hughes) and Systolic architectures for ladder forms and parallel Kalman filters (Stanford and ISI).

A brief summary of the technical work, grouped in terms of research and, is described in the following sections.

SECTION 1: SIGNAL PROCESSING ALGORITHMS AND THEORY

- Sec. 1.1: adaptive spectral estimation;
- Sec. 1.2: signal processing theory;
- Sec. 1.3: parallel signal processing algorithm;

SECTION 2: HIGHLY PARALLEL COMPUTING STRUCTURES

- Sec. 2.1: dedicated signal processing architectures;
- Sec. 2.2: array processors; and
- Sec. 2.3: VLSI implementation of signal processing architectures.

Following this summary are the progress reports from individual institutes.

1. Signal Processing Algorithms and Theory

As to the first research front, it hinges upon a thorough, in-depth understanding of mathematics and algorithm analysis. In addition to the classical mathematical techniques such as Fourier transform, linear dynamic systems, random process, etc. there arises a new signal processing mathematics branch which can be grossly termed as modern spectral analysis. Explicitly or not, a large class of signal processing applications have had extensive use of this analysis as a technical basis. Therefore, our research effort aims at developing a theoretical and algorithmic basis for modern spectral analysis methods and signal processing applications.

1.1 Adaptive Spectral Estimation Methods

1.1.1 Adaptive Notch Filtering (USC [1-3])

Using a steady state frequency domain approach, a new method has been developed for the retrieval of sinusoids/narrowband signals in additive noise colored or white. The method suggested has been shown to require smaller filter length to produce unbiased estimates, compared to the existing autoregressive method. For its implementation, a pole-zero filter where the feedback and feedforward coefficients are related (constrained ARMA), has been developed. A study of the performance and implementational aspects of the filter have been undertaken. The details of this newly developed are discussed in the full report. For a stable implementation, parallel and cascade forms have been shown to be useful. A parallel processing scheme developed shows great promise.

1.1.2 Adaptive Least Squares Ladder Form Algorithm (UCSD [17-18], ISI and Stanford [20-22])

The gradient methods of adaptive filter implementation require various adaptive power estimates to be made. The performance sensitivity of the gradient methods to different time constants of the adaptive power estimation loops has been under study at UCSD. For the specific case of underlying frequency versus time dynamics consisting of dual steps, the simulations presented in the report investigate this sensitivity in the context of a frequency tracking problem. Based on an intensive study on both the eigenvalues and singular values of the autocorrelation matrix arising from complex data, adaptive lattice structures appear to show significant performance advantages over their direct form counterparts, due to an insensitivity to eigenvalue spread. A documentation task on the eigenvalue spread present both in future controlled simulations, as well as in actual sonar data, is in progress.

In addition, a recursive least-squares ladder-form algorithm for predicted residuals rather than filtered residuals has been derived at ISI and Stanford. The reflection coefficients and the order updates of the residuals in the new algorithm are computed simultaneously. This formulation improves the throughput rate and numerical stability of existing recursive least-squares algorithms.

1.2 Signal Processing Theory

1.2.1 2-D Spectral Estimation (USC [13])

Our recent research has been concerned with developing systematic methods for 2-D spectral estimation from raw data using random field models. We assume that the given finite data is represented by an appropriate Gaussian Markov random field (MRF) model.

By using specific finite toroidal lattice representations and Gaussian maximum likelihood estimates we have developed new 2-D spectral estimates. It turns out that the MRF spectrum is also the maximum likelihood spectrum arising in frequency-wave number analysis. Furthermore, the sample correlation values of the given observations in an array N are in perfect agreement with the estimated theoretical correlations in N obtained by Fourier inverting the MRF spectrum. Thus the MRF spectrum developed by us converges to the 2-D maximum entropy spectral estimate asymptotically. Currently we have begun investigations on parallel implementation of the algorithms for 2-D spectral estimation.

1.2.2 Relationships Between Several Popular Methods for Spectral Estimation and Array Processing

It may seem too ambitious to compare all currently popular high-resolution spectral estimation methods. For example, while maximum entropy method related to autoregressive modeling is recieving a tremendous popularity, it may suffer from bias and resolution problems when additive noise is non-negligible. On the other hand, Pisarenko's method based on sinusoidal modeling enjoys relatively better performance in the presence of noise but in general suffers from numerical sensitivity problems. However, from a different perspective, Pisarenko's method can be viewed as an extension of the MEM method with the removal of the noise contribution. Therefore, an attempt is being made at developing a unified framework for the spectral analysis techniques. Moreover, the unification attempt is being extended to the counterpart of spectral analysis in array processing application. Though the covariance matrix will no longer have a Toeplitz structure and the phasing vectors are more complex in array processing situations, we are convinced that the general principles remain largely applicable. We are currently looking into theoretical and computational relevances between several modern array processing and spectrum estimation methods.

1.2.3 Toeplitz Approximation Method (USC[4])

Recently, the study on approximation theory and its applications has received considerable attention. In our work, a narrowband/sinusoidal signal retrieval problem is formulated in terms of approximation of Toeplitz autocovariance matrix. A Toeplitz approximation method based on singular value decomposition is proposed and simulation results indicate some improvement over some previously proposed methods.

1.2.4 Scattering Theory (ISI and Stanford[20])

This task has just been started and our hope is to use scattering theory to decouple large-scale, 2-D signal processing problems into blocks which can be processed simultaneously. Scattering theory will also provide physical insight into the implementation of ladder forms and parallel Kalman filters.

1.3 Parallel Algorithms

1.3.1 Parallel Algorithms for Image Processing and Analysis.(USC)

Our recent research at Image Processing Institute,USC, has been concerned with parallel algorithms for image processing and image analysis. Most of the effort has been concerned with parallel implementation of nonstationary adaptive image restoration.

Recursive and non-recursive implementation of locally adaptive restoration has been studied. These techniques estimate the local nonstationary mean and variance of ideal scenes from degraded data. Most blurring degradations are also highly local, so that local parallel processing combined with the nonstationary image model data can be used to minimize local mean-square erro (MSE) in a parallel fashion. We have shown that local MSE is not a bad error criterion for image processing, as opposed to the usual global MSE taken over the entire scene. Global MSE often does not correlate well with human observer judgments of image quality.

We have looked at the application of these techniques to systems with coherent speckle noise, such as synthetic aperature (SAR) imagery, coherent sonar and acoustic imaging. Both recursive (Kalman-like) and local sectioned parallel implementations are being studied in detail.

In addition, we have began investigations on parallel feature extraction for texture identification and texture segmentation.

1.3.2 A Parallel Algorithm for Solving Toeplitz System (USC [6-7])

We have developed a parallel algorithm for solving a Toeplitz system $Tx = y$ where T is a Toeplitz matrix, i.e., $[T]_{ij} = t_{i-j} = t_k$, $-N \leq k \leq N$. In general, solving an N by N linear system takes O(N**3) steps of operations. In contrast, the Levinson algorithm effectively utilizes the Toeplitz structure to reduce the overall computation to O(N**2) operations. The Levinson procedure, however, has to call upon an inner product operation to compute the vital reflection coefficients. In order to achieve full parallelism, we have to further explcit the Toeplitz structure. For this purpose, we have proposed a new, pipelined version of the Levinson algorithm which allows the reflection coeffecients to be computed in a pipelined fashion. This avoids the need of the inner product operations, and the total computing time is therefore reduced to O(N).

1.3.3 Toeplitz Eigenvalue Computation (USC[5])

This research task deals with the parallel computation of the minimum eigenvalue of a Toeplitz matrix. The minimum eigenvalue has an important interpretation as the power of additive, white noise to be determined in a noisy statistical environment. In many high resolution spectrum analysis problems , the estimation and removal of such noise contribution is essential for unbiased estimates. Our objective is again to derive an O(N) computation algorithm to estimate the minimum eigenvalue of a given Toeplitz covariance matrix. This goal can be accomplished by adopting the pipelined Toeplitz computing structure discussed earlier and a careful utilization of a relationship between the minimum eigenvalue and the redisues F that arise in the Levinson algorithm. Based on this relationship, a fast iterative procedure is developed to successively estimate the minimum eigenvalue. Based on simulation results for such an application, some improvements are observed in both the computing speed as well as accuracy of estimates. Although much more computational complexity analysis is yet to be demonstrated, we are convinced that this approach will have a major in future applications of high speed, high resolution spectrum estimation problems.

1.3.4 Applications of SVD to Signal Processing (NOSC, USC)

It is well known that SVD can be used in many signal processing applications. Therefore parallel (real-time) implementation has been an important research focus. Some partial results are offered in the report. The most noteworthy result is the significant numerical improvement of 60db in terms of dynamic range obtained in the computation of eigenvalue of $R = A^T A$ via SVD of A. This approach is being extended to generalized eigensystem computation.

1.3.5 Parallel Kalman Filter Algorithms (Stanford and ISI[25-27])

The research on Parallel Kalman Filters (PKF's) has been divided into two major tasks: PFK's for linear signal processing applications and PKF's for nonlinear signal processing.

For linear signal procesing applications, the predictor and corrector equations in the Kalman filter can be computed on separate processors simultaneously. A PKF has been coded and simulated. A stability analysis of the PKF is currently underway.

For nonlinear signal processing applications (such as spectrum estimation), it is difficult to decouple the extended Kalman filter (EKF) equations. Therefore, to speed-up computations parallel predictor-corrector (P^2C) methods have been used to speed-up the linearization process in the EKF. For example, the P^2C methods are 2 - 100 times faster than sequential methods given 2 - 100 processors. Applications to maximum likelihood estimation of sinusoidal signals in wide-band noise has also been studied.

1.3.6 Parallel Algorithms for Seismic Signal Processing (USC[12])

Parallel Processing techniques for generating synthetic seismograms and for the computation of the output of a horizontally stratified, non-absorptive medium propagating plane waves vertically; have been-studied.

2. Highly Parallel Computing Structures

The aforementioned research effort on signal procesing algorithm and theory, equipped with parallel algorithms, and adaptive on-line processing techniques, will serve as a useful cornerstone for real-time high performance signal processing area. However, the real major thrust for high-speed signal processing lies in effective utilization of the enormous computation capability provided by the VLSI circuits. Therefore, our other research task aims to bring the revolutionary VLSI device technology to an effective signal processing application.

2.1 Dedicated Architectures for Signal Processing

2.1.1 Pipelined Toeplitz System Solver (USC[6-7])

This new parallel algorithm for solving Toeplitz system can be implemented for parallel computation with full compliance with the VLSI communication constraint. Specifically, a pipelined processor architecture with O(N) processors is developed which uses only localized interconnections and still retains the maximum parallelism attainable.

We believe that the proposed pipelined Toeplitz system solver [5, 6] is perhaps the most efficient, fast, and practical (in VLSI sense) design available for solving Toeplitz systems. Moreover, the design methodology demonstrated in this work should also help answer some fundamental problems faced in designing of VLSI parallel processor architectures.

2.1.2 Architectures for Ladder Forms and Parallel Kalman Filters (ISI[25-27])

Ladder-form architectures for implementing the recursive least-squares ladder-form algorithm have been developed.

Existing systolic array architectures are being evaluated to determine which architectures are suitable for implementing parallel Kalman filters (PKF's). Systolic array architectures for Cholesky decomposition and triangularization have been considered to implement square-root PKF's. In addition, VLSI architectures based on mapping the Kalman filter equations directly onto the processor architectures have been examined.

2.2 Array Processors for Signal Processing

2.2.1 Wavefront Array Processor (USC[8-11])

The traditional design of parallel computers and languages is not very suitable for the design of VLSI array processors for signal processing. VLSI imposes the restrictions of local data-dependence and recursivity on the algorithms that can be handled by such an array processor. Such algorithms can be viewed as a sequence of waves (of data and computational activity). This naturally leads to a wavefront based programmable computing network, which we call the

Wavefront Array Processor (WAP).

Our contribution hinges upon the development of a wavefront-based language and architecture for a programmable special purpose multiprocessor array. Based on the notion of computational wavefront, the hardware of the processor array is designed to provide a computing medium that preserves the key properties of the wavefront. In conjunction, a wavefront language (MDFL) is introduced that drastically reduces the complexity of the description of parallel algorithms and simulates the wavefront propagation across the computing network. Together, the hardware and the language lead to a programmable Wavefront Array Processor (WAP). The WAP blends the advantages of the dedicated systolic array and the general purpose Data-Flow machine and provides a powerful tool for the high speed execution of a large class of matrix operations and related algorithms which have widespread applications.

2.2.2 Array Processor Testbed (NOSC[16])

Systolic architectures have been developed for signal processing tasks and systolic implementations have been examined for matrix multiplication, partitioned matrix inversion, computation of crossambiguity functions, formation of outer products and skewed outer products, and multiplication of arbitrary matrices by Hankel and Toeplitz matrices. Implementation of the generalized singular value decomposition of Van Loan has been identified as an important task for high resolution direction finding.

Previous systolic architectures have been modified and extended to provide improved modularity with respect to the set of functions to be performed and the size of the data arrays to be processed. Specifically, an improved systolic architecture called the "Engagement Processor", a type of wavefront processor, has been invented.

In addition, a microprocessor based 8 x 8 systolic array is being built at NOSC to serve as a test bed for future algorithm developments, architecture designs (including systolic, wavefront, and engagement processor array) and VLSI implementation.

2.3 VLSI Implementation (Hughes[19])

A Hardware Implementations of the Toeplitz System Solver and other related systems e.g. for Fadeev algorithm, in VLSI was undertaken at Hughes (HRL(19))

Our work has included an extensive investigation of the various possible hardware implementations (fixed point versus floating point, serial/parallel versus parallel only, etc.) and arithmetic algorithms. Chip organization, pin-out problems, cell designs, and chip-to-chip communications are also considered.

There are several important design and architectural considerations that make our approach well suited to the capabilities of VLSI technologies. These are summarized below:

- Identical processing elements
- Local communication
- Expandability, allowing chips to function on arbitrary-sized kernals
- Local data storage
- exploitation of full concurrency (minimum number of idle processors).

3. SUMMARY

In conclusion, in order to keep pace with the rapid advance in the VLSI technology, the signal processing community should not only look into advanced processing theory and parallel processing methods, but also exercise a timely influence on architectural design of future VLSI computing structures. This is exactly the goal of our ONR-SRO II project. Keeping up the current momentum , our joint effort will definitely represent a major contribution to the modern VLSI Signal processing technology.

4. Summary of report - UCSD

MPL-U-9/82

ADAPTIVE LEAST SQUARES LATTICE STRUCTUES AND THEIR APPLICATIONS TO PROFLEMS IN UNDERWATER ACOUSTICS

Progress Report

1 September 1981 - 28 February 1982

W. S. Hodgkiss

I. BACKGROUND

Numerous applications exist which require a linear filtering operation. Often, the nature of that filtering task (e.g. spectral shaping characteristics) is time varying in some nondeterministic fashion due to nonstationarity of the underlying time series. In such situations, a filter which can adapt to a changing environment is needed. For the purpose of derivation, the adaptive filter must have both a well-defined goal (e.g. linear prediction) and a well-defined performance measure (e.g. weighted summation of the squared prediction error residuals). Once derived, it is extremely important to understand the performance characteristics of the adaptive filter both from a theoretical, as well as a practical point of view prior to its use in the processing or real data.

II PROGRESS: 1 Septemer 1981 - 28 February 1982

The gradient methods of adaptive filter implementation require various adaptive power estimates to be made. The performance sensitivity of the gradient methods to different time constants of the adaptive power estimation loops has been under study. For the specific case of underlying frequency versus time dynamics consisting of dual steps, the simulations presented in [1] and [2] investigate this sensitivity in the context of a frequency tracking problem.

In addition to these performance studies, a substantial amount of time has been devoted towards implementing software for the calculation of both the eigenvalues and singular values of the autocorrelation matrix arising from complex data. Adaptive lattice structures appear to show significant performance advantages over their direct form counterparts due to an insensitivity to eigenvalue spread. Our desire is to be able to document the eigenvalue spread present both in future controlled simulations, as well as in actual sonar data which we will be processing.

5. Summary of Report - HRL

Concurrent VLSI Architectures for Matrix Operations and Linear Systems Solution(J.G. Nash, G.R. Nudd, S. Hansen)

The goal of this project is to develop novel, special purpose computing structures capable of throughputs far exceeding that available from present day commercial computers. To achieve this goal we are exploiting the concurrency potentially available in algorithms associated with a wide range of applications (automated production techniques including image analysis, robotics control and solution of previously intractable simulation problems). The availability of Very Large Scale Integration (VLSI) has made this approach economically feasible and at the same time has considerably enhanced circuit performance.

Matrix operations represent a major part of the computational requirement encountered in many computer applications. Examples of matrix operations are multiplication, inversion, and L-U decomposition. In signal processing such operations can be found in adaptive filtering, data compression, cross-ambiguity calculations, and beamforming. In robotics a matrix formulation of the control of manipulator joints in a Cartesian coordinate system is the most straightforward and convenient. Image analysis and restoration is often based on "kernal" operations over a relatively small window of pixels, e.g., 30 x 30, as might be found in relaxation techniques. Therefore, specification of image processing algorithms in terms of matrix operations occurs in a natural way.

Matrix operations are well-suited to concurrent implementations in which a number of small, identical processors operate simultaneoulsy on the matrix elements. Thus, a set of "matrix operator" chips, made using VLSI, would, when coupled to a general purpose host computer, provide both a high computational throughput and the flexibility to perform a wide range of algorithms spanning many applications. This is a far more efficient approach than mapping a particular algorithm directly into hardware.

The increase in throughput obtainable by using an array of processors operating concurrently is proportional to the number of processors. Thus, a matrix inversion, which takes $O(n^3)$ steps can be performed in $O(n)$ time steps using an n x n array of processors. The speedup in time is then a factor of n^2.

As an example of the utility of the matrix approach, consider the requirements for control of servoing typical manipulators with approximately 1 meter of reach. The most time consuming part of this computation is the calculation of the pseudo inverse of the Jacobian. A matrix formulation of the control problem would require approximately 2,000 multiplications in 250 microseconds.[1] Assuming that one multiplication could be performed every microsecond, a serial processor would take 2 milliseconds; however, a 6 x 6 array of processors, each with the same 1 microsecond multiplication time, would complete the calculation in approximately 55 microseconds.

The processor arrays we are investigating incorporate important design and architectural considerations that make our approach well suited to the capabilities of VLSI technologies. These are:

- Identical processing elements
- Local communication
- Expandability, allowing chips to function on arbitrary-sized kernals
- Local data storage
- Exploitation of full concurrency (minimum number of idle processors).

Each processing element consists of a basic inner-product ("ax+b") calculator and communicates only with its nearest North, South, West, East, and diagonal neighbor.

We will describe in this report two basic systems which we have been investigating. The Toeplitz linear system solver, first suggested by S.Y. Kung,[2] is based on the Weiner-Levinson algorithm and is suitable for operations on a Toeplitz-type matrix (elements along any diagonal are identical). Toeplitz matrices appear in stationary signal processing applications. For example, the autocorrelation matrix is often of this type. The second system we are investigating is based on the Faddeev algorithm and is suited for general matrices where there is no overriding symmetry. In Section 2 we describe these algorithms and their functional architectural embodiments.

In Section 3 we report on numerous detailed numerical studies we have performed. These statistical studies were done by simulating the various architectures using the high level programming language, APL, which has the capability of operating on arrays of numbers in a straightforward and intuitive way. The simulations have served three main functions. First, they have provided a means for verifying the correctness of the data flow within the Toeplitz and Faddeev matrix processors. Second, they have allowed us to study the numerical stability of the algorithms for representative signal processing data. And third, they have been useful in predicting effects of finite register lengths, roundoff techniques, and pivoting schemes.

In Section 4 we describe the hardware implementation of the Toeplitz Linear System Solver. This required an extensive investigation of the various possible hardware implementations (fixed point versus floating point, serial/parallel versus parallel only, etc.) and arithmetic algorithms. Chip organization, pin-out problems, cell designs, and chip-to-chip communications are also discussed and the final chip layout shown.

6. Summary of Report - NOSC

Signal Processing With Systolic Arrays

J. M. Speiser and H. J. Whitehouse

Naval Ocean Systems Center
San Diego, CA. 92152

OBJECTIVE: To devise parallel computing architectures for present and future Navy needs in real-time signal processing, with emphasis on spectrum analysis and direction finding.

APPROACH: Select the needed signal processing operations which present a large computational load, and determine corresponding computational algorithms which exhibit parallelism and regularity of data flow. Structure such algorithms as systolic architectures and examine the required arithmetic operations and I/O between successive algorithms or systolic subsystems.

PROGRESS: (through Feb 82) systolic architectures have been developed for signal processing tasks and systolic implementations have been examined in coordination with the NOSC-university-industry SRO task team on spectrum analysis: partitioned matrix multiplication, partitioned matrix inversion, computation of crossambiguity functions, formation of outer products and skewed outer products, and multiplication of arbitrary matrices by Hankel and Toeplitz matrices. Implementation of the generalized singular value decomposition of Van Loan has been identified as an important task for high resolution direction finding. Two sessions on parallel processing algorithms and architectures were organized for the SPIE International Symposium in San Diego, August, 1981.

Previous systolic architectures have been modified and extended to provide improved modularity with respect to the set of functions to be performed and the size of the data arrays to be processed. Specifically, an improved systolic architecture called the "Engagement Processor", a type of wavefront processor, has been invented and has the following advantageous attributes:

a) Permits more efficient multiplication of dense matrices and explicitly provides for partitioned multiplication of matrices having more elements than the number of processors in the array. Efficiency is approximately 1/3 for a single matrix multplication of matrices which match the processor in size, and nearly 1 for the pipelined multiplication of successive matrices, including the partitioned multiplication of large matrices.

b) Can include the modified hexagonal array of H. T. Kung as a subset permitting L-U decomposition of matrices using the same processor.

c) Permits efficient utilization of the processors for covariance estimation for array data.

d) Permits calculation of an N point DFT in about 11 N.5 arithmetic cycles using N processors, including I/O time for multiplexed I/O.

For matrix multiplication, covariance estimation, and DFT calculation, the engagement processor may be viewed as a rectangular systolic array with two types of memory augmentation: local memories for the inner product step processors, permitting parallel access to stored arrays, and two sets of virtual delay lines, acting as I/O buffers at two edges of the array. The same processor array may be used to perform L-U decomposition of matrices with only minor modification because of the known result of S. Y. Kung that a hexagonal systolic array may be viewed as a rectangular array with diagonal interconnects.

The engagement processor systolic architecture has been extended to provide faster matrix multiplication with only a nominal increase in hardware complexity. One extension uses a single "bus expander" to permit multiplication of an arbitrary real N by N matrix with a Toeplitz of Hankel matrix with only vector storage. The signal processing operations of outer product formation triple product convolution and skewed outer product formation were also examined for timing on an engagement processor. Combinations of these and the previously described operations were examined in order to compare four algorithms for crossambiguity function calculation on a systolic processor - the algorithms previously examined for optical processor implementation: a two-dimensional Fourier transform technique, the combination of a skewed outer product and one-dimensional DFT (as used in a τ-slice processor), the combination of pointwise multiplication and Hankel matrix multiplication (as used in a space-integrating processor), and a simulated triple product convolution (as used in a time-integrating processor). The last two techniques were significantly faster on an engagement processor than the first two: time 4N versus 13N and 7N. Also, a technique was devised to combine toroidal interconnection of an engagement processor with a set of N bus expanders to permit the multiplication of arbitrary N by N matrices in time N using N^2 processors. This corresponds to 100% efficient use of the processors without requiring pipelining of successive matrix multiplications, but requires non-nearest neighbor interconnections.

We are currently examining parallel algorithms and systolic architectures for orthogonal-triangular factorization of matrices for least squares solution and eigensystems via the QR algorithm, as well as constrained least squares solution and recursive updating of least squares solutions. Current bottlenecks appear to be accumulation of the orthogonal matrix, incorporation of shifts, and the incorporation of pivoting in a partitioned QR decomposition.

We have observed that in most signal processing applications where the eigensystem of a covariance matrix is desired, the matrix is estimated as $\hat{R} = A^T A$, where A is experimentally observed. It is therefore numerically desirable to find the eigensystem of $\hat{R}$ from the singular value decomposition of A rather than computing directly with $\hat{R}$. If the SVD of A is $A = PDQ^T$, where P and Q are orthogonal matrices and D is diagonal, then $\hat{R} = A^T A = Q^T D^2 Q$, so the eigenvalues of $\hat{R}$ may be computed as the squares of the singular values of A, and the eigenvectors of $\hat{R}$ are the right singular vectors of A. This suggestion has been applied by the Marine Physical Laboratory to the estimation of the spectrum of a multiple tone signal. Using PDP-11 floating point arithmetic, the computational dynamic range was increased by approximately 60 dB when the eigenvalues of $\hat{R}$ were computed as the squares of the singular values of A. We are currently examining the applicability of generalized eigensystens to Navy direction finding problems. We propose to extend this approach to the generalized eigensystem computation, by using the generalized singular value decomposition of Van Loan.

As alternatives to the QR algorithm for the eigensystem problem and the Golub adaptation of QR to the singular value decomposition, we are examining Jacobi methods for the eigensystem problem and the one-sided orthogonalization adaptions of the Jacobi method for the singular value decomposition. Although the Jacobi and related methods require approximately three times the number of multiplication of the corresponding QR methods, the Jacobi methods have two strong advantages: a) easier parallelization; b) in many signal processing applications a great speedup is possible by applying a preconditioning transformation.

For problems requiring the eigensystems of the covariance matrix of a wide-sense stationary random process, or the corresponding singular value decomposition of A, where $\hat{R} = A^T A$, the basis vectors of the discrete Fourier transform are an approximate eigensystem for $\hat{R}$, and may be used to approximately diagonalize $\hat{R}$ by a unitary similarity transformation, or to perform an approximate one-sided orthogonalization on A. The corresponding Jacobi iterations or one-sided rotations may be used to rapidly improve the decomposition.

Publications during this period have included a tutorial on recent parallel architectures [10], an exposition of the application of current systolic architectures to sonar problems [11], and a description of the NOSC systolic testbed for architecture validation and algorithm development [12].

In order to facilitate the development of algorithms and software for wavefront processors, NOSC will provide an 8 x 8 programmable systolic array[12] to USC.

PUBLICATIONS:

1. Speiser, J.M. and H.J. Whitehouse, Architectures for Real-Time Matrix Operations, to appear in proceedings of GOMAC-80, Government Microcircuit Applications Conference, Houston, TX, 19-21 Nov 1980.

2. Speiser, J.M., H.J. Whitehouse, and K. Bromley, Signal Processing Applications for Systolic Arrays, to appear in proceedings of the 14th Asilomar Conference on Circuits, Systems, and Computers, held at Pacific Grove, CA, 17-19 Nov 1980.

Aczel, J., J.A. Baker, K. Davidson, B. Forte, and F. Neuman, Application of Functional Equations to Factorization of Transform Kernels for Use in Signal Processing, University of Waterloo progress report on U.S. Navy Contract N66001-80-C-0518, Dec. 1980.

4. Priester, R.W., K. Bromley, H.J. Whitehouse, and J.B. Clary, Signal Processing with Systolic Arrays, to appear in proceeding Agard Symposium on Tactical Airborne Distributed Computing Networks, held at ROROS, Norwary, 22-25 Jun 1981.

5. Bromley, K. and H.J. Whitehouse, Signal Processing Technology Overview, SPIE Vol. 298, Real Time Signal Processing IV, paper 298-14.

6. Speiser, J.M. and H.J. Whitehouse, Parallel Processing Algorithms and Architectures for Real-Time Signal Processing, presented at the SPIE International Technical Symposium, San Diego, CA., 25-28 Aug 1981. To appear in SPIE Vol. 298, "Real-Time Signal Processing IV", paper 298-01.

7. Bromley, K., J.J. Symanski, J.M. Speiser, and H.J. Whitehouse, Systolic Array Processor Developments, to appear in proceedings of the CMU Conference on VLSI Systems and Computation, Carnegie Mellon University, Pittsburgh, PA., 19-21 Oct 1981.

8. Speiser, J.M., K.Bromley, and H.J. Whitehouse, Signal Processing Applications of Real-Time Matrix Operations, in the proceedings of the La Jolla Institute Workshop on Research Requirements for Advanced Computer Design, held at La Jolla, CA.

9. Priester, R., J. Clary, K. Bromley, and H.J. Whitehouse, Problem Adaptation to Systolic Arrays, to appear in SPIE Vol. 298, Real-Time Signal Processing IV, paper 298-05. Presented at the SPIE International Technical Symposium, San Diego, CA., 25-28 Aug 1981.

10. Bromley, K. and H.J. Whitehouse, "Signal Processing Technology Overview", paper 298-14 in SPIE Vol. 298, Real-Time Signal Processing IV, presented at the SPIE International Symposium, San Diego, CA, Aug 25, 1981.

11. Whitehouse, H.J. and J.M. Speiser, "Sonar Applications of Systolic Array Technology", to appear in the IEEE EASCON Conference Record, Washington D.C., Nov 17-19, 1981

12. Symanski, J.J., "A Systolic Array Processor Implementation", paper 298-14 in SPIE Vol. 298, Real-Time Signal Processing V, presented at the SPIE International Symposium, San Diego, CA Aug 25, 1981.

7. SUMMARY OF REPORT ISI AND STANFORD

PARALLEL ALGORITHMS AND ARCHITECTURES FOR NONLINEAR SIGNAL PROCESSING
(R.H. Travassos)

TABLE OF CONTENTS

SECTION 1

PROJECT STATUS

This report describes the research activities of Integrated Systems, Inc. (ISI) on contract number N00014-81-X-0191 for the period from 1 October 1981 to 1 January 1982. The research activities have focussed on five major tasks (see Figure 1). Project spending is indicated in Figure 2. The proposed research on nonlinear signal processing algorithms and architectures is being performed on schedule and within budget. A summary of the technical work is described in Section 1.1 - 1.5 and in the attached technical memos.

1.1 LADDER-FORM ALGORITHMS

A recursive least-squares ladder-form algorithm for predicted residuals rather than filtered residuals has been derived. The reflection coefficients and the order updates of the residuals in the new algorithm are computed simulateously. This formulation improves the throughput rate and numerical stability of existing recursive least-squares algorithms. The details of the newly developed ladder-form algorithm are discussed in ISI Technical Memo 5016-03.

1.2 PARALLEL KALMAN FILTER ALGORITHMS

The research on Parallel Kalman Filters (PKFs) has been divided into two major tasks: PKFs for linear signal processing applications and PKFs for nonlinear signal processing.

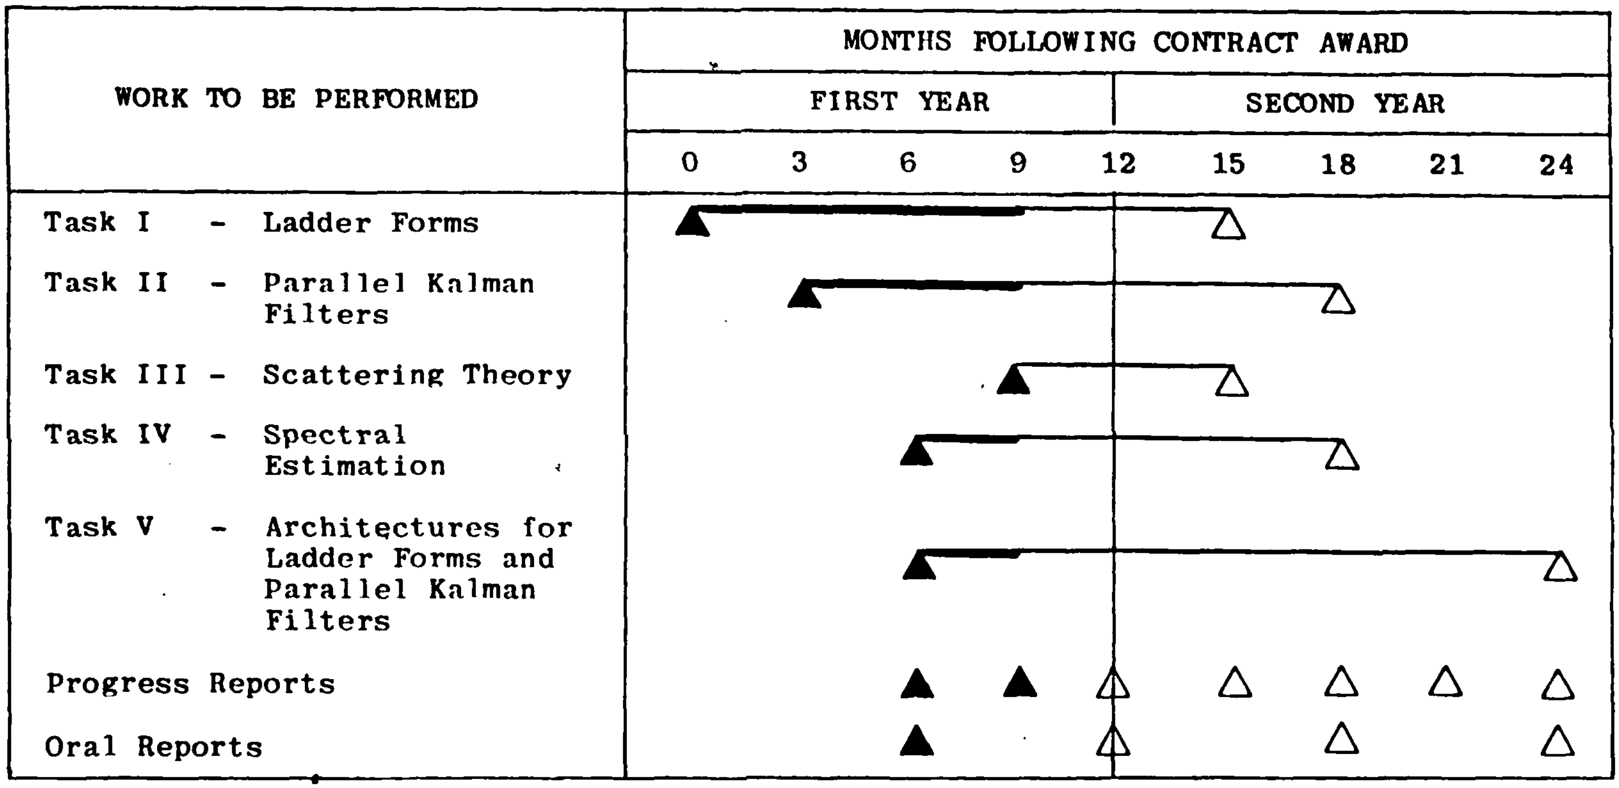

FIGURE 2. HIGHLY PARALLEL MODERN SIGNAL PROCESSING SCHEDULE

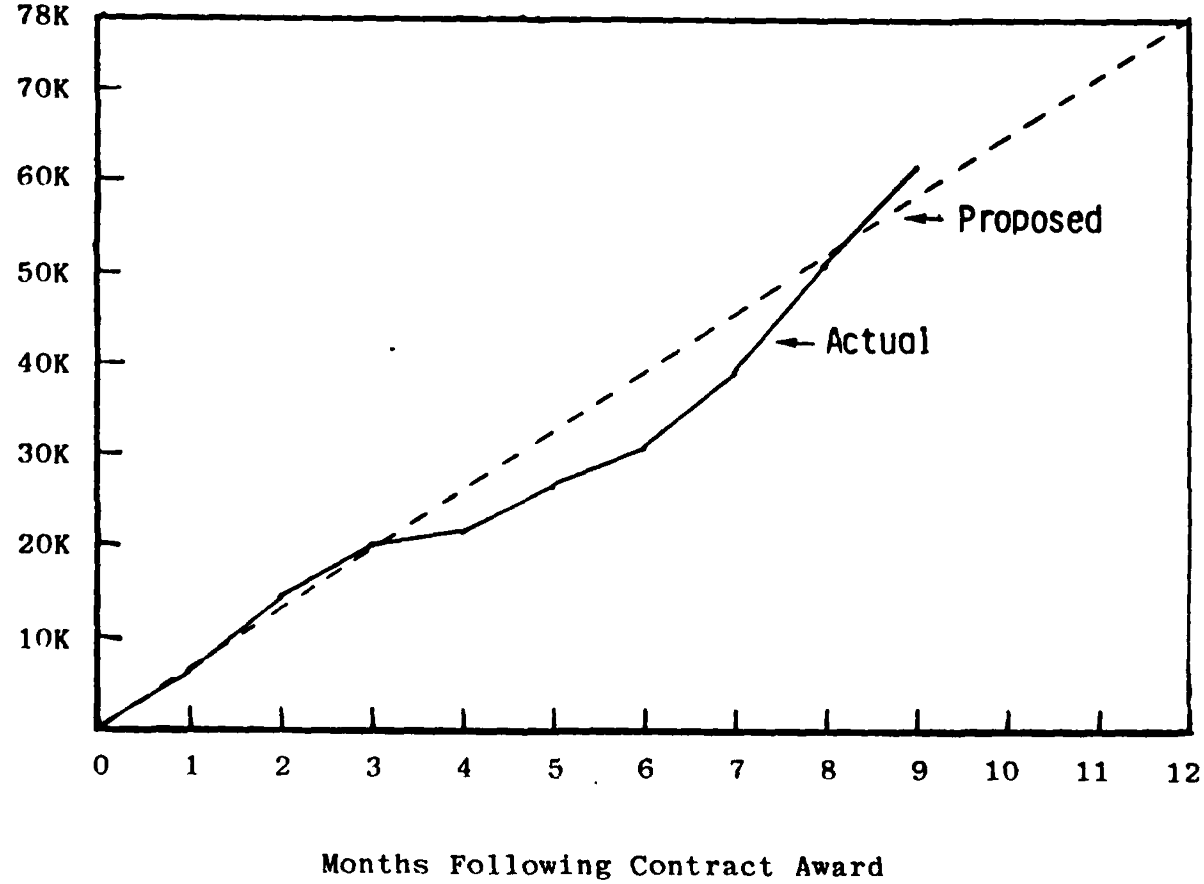

Figure 2 First Year Project Spending

1.2.1 PKFs for Linear Signal Processing Applications

For linear signal processing applications, the predictor and corrector equations in the Kalman filter can be computed on separate processors simultaneously. A PKF has been coded and simulated. A stability analysis of the PKF is currently underway. Systolic array architectures for implementing the PKF with VLSI technology have been investigated (see Section 1.5).

1.2.2 PKFs for Nonlinear Signal Processing Applications

For nonlinear signal processing applications (such as spectrum estimation), it is difficult to decouple the extended Kalman filter (EKF) equations. Therefore, to speed-up computations parallel predictor-corrector (P^2C) methods have been used to speed-up the linearization process in the EKF. The P^2C methods are 2 - 100 times faster than sequential methods given 2 - 100 processors. Research on methods for varying the integration step size has continued to make the P^2C methods more efficient.

The parallel EKF has been applied to maximum likelihood estimation of sinusoidal signals in wide-band noise. Square-Root-Free Parallel Quasi-Newton methods have been used to improve the maximum likelihood parameter estimates (see ISI Technical Memo 5016-04).

1.3 SCATTERING THEORY

This task has just been started (see Figure 1). Our hope is to use scattering theory to decouple large-scale, 2-D signal processing problems into blocks which can be processed simultaneously. Scattering theory will also provide physical insight into the implementation of ladder forms and parallel Kalman filters.

1.4 SPECTRAL ESTIMATION

A prediction error approach based on auto-regressive-moving-average (ARMA) models and a maximum likelihood (ML) method which utilizes the PKF developed under Task 2 have been coded and applied to power spectrum estimation (see ISI Technical Memos 5016-04 and 5016-05). The results indicated that: (1) ARMA based methods gave sharp peaks compared with auto-regressive (AR) techniques and (2) the PKF gave excellent ML estimates of sinusoidal parameters even under poor signal-to-noise-ratio conditions. Therefore, research on ARMA and ML based spectral estimation will continue.

1.5 ARCHITECTURES FOR LADDER FORMS AND PARALLEL KALMAN FILTERS

Ladder-form architectures for implementing the recursive least-squares ladder-form algorithm developed under Task 1 are described in ISI Technical Memo 5016-03.

Existing systolic array architectures are being evaluated to determine which architectures are suitable for implementing parallel Kalman filters (PKFs). Systolic array architectures for Cholesky decomposition and triangularization have been considered to implement square-root PKFs. In addition, VLSI architectures based on mapping the Kalman filter equations directly onto the processor architectures have been examined. The results of this study will be included in the final report.

SECTION 2

FUTURE WORK

Future research will be conducted in accordance with the program schedule (see Figure 1). Research will continue on ladder forms, parallel Kalman filters, and spectral estimation. Emphasis will be placed on: (1) parallel architectures which are suitable for implementation with VLSI/VHSIC technology, (2) stability analysis of the PKF algorithms, (3) extending the PKF algorithms and architectures for 2-D signal processing applications (e.g., image processing), and (4) integrating the research activities of the SRO participants (e.g., using the data flow language to emulate the PKF architectures). Simulation studies will continue to evaluate the performance of the newly developed parallel algorithms and architectures for spectral estimation.

SECTION 3

TECHNICAL PUBLICATIONS

The following technical papers have been written under contract number N00014-81-K-0191:

1. Travassos, R. H., "Parallel Processing Algorithms for Unconstrained Minimization," prepared for the Journal of Optimization Theory and Application, July 1981.

2. Travassos, R. H., "Square-Root Parallel Quasi-Newton Methods for Nonlinear Optimization," invited paper presented in the special session on parallel optimization techniques, Optimization Days, Montreal, Canada, May 1981.

3. Travassos, R. H., "Parallel Processing Algorithms for System Parameter Identification," prepared for the IFAC Symposium on System Identification, Washington, D. C., June 1982.

4. Travassos, R. H., "Parallel Kalman Filtering," ISI Technical Memo 5016-01, October 1981.

5. Jover, J. M. and Travassos, R. H., "Derivation of Algorithms for UDU^T Kalman Filters," ISI Technical Memo 5016-02, October 1981.

6. Reddy, V. U., "A Numerically Stable and High-Speed Recursive Least-Squares Ladder-Form Algorithm," ISI Technical Memo 5016-03, November 1981.

7. Travassos, R. H., "Maximum Likelihood Estimation of Sinusoidal Signals in Wide-Band Noise," ISI Technical Memo 5016-04, January 1982.

8. Reddy, V. U., Travassos, R. H., and Kailath, T., "A Comparison of Nonlinear Spectral Estimation Techniques," ISI Technical Memo 5016-05, January 1982.

9. Travassos, R. H. and Andrews, A., "VLSI Implementation of Parallel Kalman Filters," prepared for the AIAA Guidance and Control Conference, San Diego, California, August 1982.

8. SUMMARY OF REPORT - USC

Table of Contents

1 PARALLEL COMPUTING STRUCTURES (S.Y.Kung, Y.H.Hu, R.Gal-ezer, K.S.Arun ,D.V.B.Rao)

With the rapidly growing microelectronics technology leading the way, modern signal processor architectures are undergoing a major revolution. The availability of low cost, fast VLSI (Very Large Scale Integration) devices promises the practice of cost-effective, high speed, parallel processing of large volume of data. This makes possible ultra high throughput-rate and therefore, designates a major technological breakthrough for real-time signal processing applications. On the other hand, it has become more critical than ever to gain a fundamental understanding of the algorithm structure, architecture, and implementation constraints in order to realize the full potential of VLSI computing power. In our work, the two most critical issues - parallel computing algorithm and VLSI architectural constraint will be considered:

1. To structure the algorithm to achieve the maximum parallelism and, therefore, the maximum throughput-rate.

2. To cope with the communication constraint so as to compromise least in processing throughput-rate.

1.1 A highly concurrent Toeplitz system solver [5-6]

Based on the above considerations, we have developed a highly concurrent Toeplitz system solver, featuring maximum parallelism and localized communication.

Toeplitz systems arise in numerous, wide-spread applications ranging from speech, image, neurophysics, to radar, sonar, geophysics, and astronomical signal processing . Our contribution lies in the development of a highly concurrent algorithm and pipelined architecture which is able to solve a Toeplitz system in O(N) processing time in an array proecssor, as opposed to O(N**3) for general (sequential) Gauss elimination procedure or O(N**2) for (sequential) Levinson algorithm.

For parallel consideration, we note that the Levinson procedure has to call upon an inner product operation to compute the vital "reflection coeffecients". Even when N processors is utilized, an inner product operation will need at least logN units of time. This will amount to a total of O(NlogN) units of computing time for the entire Levinson procedure. This is of course unsatisfactory since the processors are not effectively utilized.

In order to achieve full parallelism, we have to further exploit the Toeplitz structure. For this purpose, we have proposed a new, pipelined version of the Levinson algorithm which allows the reflection coefficients to be computed in a pipelined fashion. This avoids the need of the inner product operations, and the total computing time is therefore reduced to O(N).

This new algorithm can be implemented in full compliance with the VLSI communication constraint. More precisely, a pipelined processor architecture is developed which uses only localized interconnections and still retains the maximum parallelism attainable.

In summary, we believe that the proposed pipelined Toeplits system solver is perhaps the most efficient, fast, and practical (in VLSI sense) design available for solving Toeplitz systems. Moreover, the design methodology demonstrated in this work should also help answer some fundamental problems faced in designing of VLSI parallel processor architectures.

1.2 Toeplitz Eigenvalue Computation [36]

This research task deals with the parallel computation of the minimum eigenvalue of a Toeplitz matrix. The minimum eigenvalue has an important interpretation as the power of additive, white noise to be determined in a noisy statistical environment. In many high resolution spectrum analysis problems, the estimation and removal of such noise contribution is essential for unbiased estimates. Our objective is again to derive an O(N) computation algorithm to estimate the minimum eigenvalue of a given Toeplitz covariance matrix. This goal can be accomplished by adopting the pipelined Toeplitz computing structure discussed earlier and a careful utilization of a relationship between the minimum eigenvalue and the redisues E that arise in the Levinson algorithm. Based on this relationship, a fast iterative procedure is developed to successively estimate the minimum eigenvalue. Based on simulation results for such an application, some improvements are observed in both the computing speed as well as accuracy of estimates. Although much more computational complexity analysis is yet to be demonstrated, we are convinced that this approach will have a major in future applications of high speed, high resolution spectrum estimation problems.

1.3 Wavefront Array Processor

The traditional design of parallel computers and languages usually suffers from heavy supervisory overhead incurred by synchronization, communication, and scheduling tasks, which severely hamper the throughput rate which is critical to real-time signal processing.

Furthermore, additional restrictions imposed by VLSI will render the general purpose array processor very inefficient. We therefore restrict ourselves to a special class of applications, i.e. recursive and local data dependent algorithms, to conform with the constraints imposed by VLSI. however, this restriction incurs little loss of generality, as a great majority of signal processing algorithms possess these properties. One typical example is a class of matrix algorithms.

Very significantly, these algorithms involve repeated application of relatively simple operations with regular localized data flow in a homogeneous computing network. This leads to an important notion of computational wavefront, which portrays the computation activities in a manner resembling a wave propagation phenomenon. More precisely, the recursive nature of the algorithm, in conjunction with the localized data dependency, points to a continuously advancing wave of data and computational activity.

The wavefront concept, provides a firm theoretical foundation for the design of highly parallel array processors and concurrent languages. Moreover, this concept appears to have some distinct advantages.

Firstly, the wavefront notion drastically reduces the complexity in the description of parallel algorithms. The mechanism provided for this description is a special purpose, wavefront-oriented language. Rather than requiring a program for each processor in the array, this language allows the programmer to address an entire front of processors.

Secondly, the wavefront notion leads to a wavefront- based architecture that conforms with the constraints of VLSI, and supports a major class of signal processing algorithms. As a consequence of Huygen's principle, wavefronts should never intersect. With a wavefront architecture that provides asynchronous waiting capability, this principle is preserved. Therefore, the wavefront approach can cope with timing uncertainties, such as local clocking, random delay in communications and fluctuations of computing-times. In short, there is no need for global synchronization.

Thirdly, the wavefront notion is applicable to all VLSI signal processing algorithms that possess locality and recursivity, and hence, has numerous applications.

The integration of the wavefront concept, the wavefront language and the wavefront architecture leads to a programmable computing network, which we will call the WAVEFRONT ARRAY PROCESSOR (WAP).The WAP is,in a sense,an optimal tradeoff between the globally synchronized and dedicated systolic array (that works on a similar set of algorithms), and the general-purpose data-flow multiprocessors. It provides a powerful tool for the high speed execution of a large class of algorithms which have widespread applications. The applications are very broad including PDE solver, SVD, linear systems solvers, sorting and searching routines.

There exist two approaches approaches to programming the WAP: a local approach, describing the actions of each processing element, and a global approach, describing the actions of each wavefront. To allow the user to program the WAP in both these fashions, two versions of MDFL are proposed: global and local MDFL. A global MDFL program describes the algorithm from the view-point of a wavefront, while a local MDFL program describes the operations of an individual processor. More precisely, the perspective of a global MDFL programmer is of one wavefront passing across all the processors, while the perspective of a local MDFL programmer is that of one processor encountering a series of wavefronts.

In summary, our contribution hinges upon the development of a wavefront-based language and architecture for a programmable special purpose multiprocessor array. Based on the notion of computational wavefront, the hardware of the processor array is designed to provide a computing medium that preserves the key properties of the wavefront. In conjunction, a wavefront language (MDFL) is introduced that drastically reduces the complexity of the description of parallel algorithms and simulates the wavefront propagation across the computing network. Together, the hardware and the language lead to a programmable Wavefront Array Processor (WAP). The WAP blends the advantages of the dedicated Systolic array and the general purpose Data-Flow machine and provides a powerful tool for the high speed execution of a large class of matrix operations and related algorithms which have widespread applications.

2 SIGNAL PROCESSING ALGORITHMS AND THEORY (S.Y.Kung, Y.H.Hu, D.V.B.Rao)

As to this research front, it hinges upon a thorough, in-depth understanding of mathematics and algorithm analysis. In addition to the classical mathematical techniques such as Fourier transform, linear dynamic systems, random process, etc. there arises a new signal processing mathematics branch which can be grossly termed as modern spectral analysis. Explicitly or not, a large class of signal processing applications have had extensive use of this analysis as a technical basis. Therefore, our research effort aims at developing a theoretical and algorithmic basis for modern spectral analysis methods and signal processing applications.

2.1 Adaptive Notch Filtering (USC [1-3])

Using a steady state frequency domain approach, a new method has been developed for the retrieval of sinusoids/narrowband signals in additive noise colored or white. The method suggested has been shown to require smaller filter length to produce unbiased estimates, compared to the existing autoregressive method. For its implementation, a pole-zero filter where the feedback and feedforward coefficients are related (constrained ARMA), has been developed. A study of the performance and implementational aspects of the filter have been undertaken. The details of this newly developed are discussed in the full report. For a stable implementation, parallel and cascade forms have been shown to be useful. A parallel processing scheme developed shows great promise.

2.2 Relationships Between Several Popular Methods for Spectral Estimation and Array Processing

It may seem too ambitious to compare all currently popular high-resolution spectral estimation methods. However, from a different perspective, Pisarenko's method can be viewed as an extension of the MEM method with the removal of the noise contribution. Therefore, an attempt is being made at developing a unified framework for the spectral analysis techniques. Moreover, the unification attempt is being extended to the counterpart of spectral analysis in array processing application, for which we are convinced that the general principles remain largely applicable. We are currently looking into theoretical and computational relevances between several recent array processing and spectrum estimation methods.

2.3 Toeplitz Approximation Method (USC[4])

Recently, the study on approximation theory and its applications has received considerable attention. In our work, a narrowband/sinusoidal signal retrieval problem is formulated in terms of approximation of Toeplitz autocovariance matrix. A Toeplitz approximation method based on singular value decomposition is proposed and simulation results indicate some improvement over some previously proposed methods.

3 REVIEW OF RESEARCH ACTIVITIES IN IPI, USC (A.A.Sawchuk, R.Chellappa)

3.1 Parallel Algorithms for Image Processing and Analysis

Our recent research at Image Processing Institute,USC, has been concerned with parallel algorithms for image processing and image analysis. Most of the effort has been concerned with parallel implementation of nonstationary adaptive image restoration. Recursive and non-recursive implementation of locally adaptive restoration has been studied. These techniques estimate the local nonstationary mean and variance of ideal scenes from degraded data. Most blurring degradations are also highly local, so that local parallel processing combined with the nonstationary image model data can be used to minimize local mean-square erro (MSE) in a parallel fashion. We have shown that local MSE is not a bad error criterion for image processing, as opposed to the usual global MSE taken over the entire scene. Global MSE often does not correlate well with human observer judgments of image quality.

We have looked at the application of these techniques to systems with coherent speckle noise, such as synthetic aperature (SAR) imagery, coherent sonar and acoustic imaging. Both recursive (Kalman-like) and local sectioned parallel implementations are being studied in detail. In addition, we have began investigations on parallel feature extraction for texture identification and texture segmentation.

3.2 Two Dimensional Spectral Estimation

Two-dimensional spectral estimation is of interest in image restoration, filtering of SAR images and texture classification. Our recent research has been concerned with developing systematic methods for 2-D spectral estimation from raw data using random field models. We assume that the given finite data is represented by an appropriate

Gaussian Markov random field (MRF) model.

This assumption reduces the spectral estimation problem to that of estimating the appropriate structure and the parameters of the model. By using specific finite toroidal lattice representations and Gaussian maximum likelihood estimates we have developed new 2-D spectral estimates. It turns out that the MRF spectrum is also the maximum likelihood spectrum arising in frequency-wave number analysis. Furthermore, the sample correlation values of the given observations in an array N are in perfect agreement with the estimated theoretical correlations in N obtained by Fourier inverting the MRF spectrum. Thus the MRF spectrum developed by us converges to the 2-D maximum entropy spectral estimate asymptotically. Currently we have begun investigations on parallel implementation of the algorithms for 2-D spectral estimation.

In addition, we are also investigating the use of another class of random field models known as spatial autoregressive models which are white noise driven non causal models for spectral estimation.

4 PARALLEL PROCESSING TECHNIQUES FOR SEISMIC PROCESSING (J.Mendel, J.Goutsias)

Because of the large volume of information involved in the simulation and processing of seismic data, and the amount of processing required, parallel techniques have begun to be studied. The recent development of VLSI systems and the growing sophestication in the design of array processors xan lead to the effecient simulation of large seismic models. We are examining some possible parallel processing techniques for the computation of the output of a horizontally stratified, non absorbtive medium in which there are vertically travelling plane compressional waves.

This task has just been started and we intend to look at different parallel structures for generating synthetic seismograms.

Selected publications under Contract number ONR N00014-81-K-0191 :

University Of Southern California

1. S.Y.KUNG and D.V.Bhaskar Rao, " An Unbiased Adaptive Method For Retrieval Of Sinusoidal Signals in Colored Noise", Proc. IEEE 20th CDC, San Diego, CA, Dec 1981

2. S.Y.Kung and D.V.Bhaskar Rao, " New Unbiased Methods For Narrowband Spectral Estimation", IFAC, NEW DELHI, India, 1982

3. S.Y.Kung and D.V.Bhaskar Rao, " Analysis and Implementation Of The Adaptive Notch Filter For Frequency Estimation", Proc. ICASSP, PARIS, MAY 3-5, 1982.

4. S.Y.Kung, " A Toeplitz Approximation Method and Some Applications", Int. Symp. on Math. Soc. of networks and Systems, Santa Monica, CA, Aug. 5-7, 1981

5. Yu-Hen Hu and S.Y.Kung, "Computation Of Minimum Eigenvalue Of Toeplitz Matrix By Levinson Algorithm ", 25th SPIE Conf., San Diego, 1981.

6. S.Y.Kung and Yu-Hen Hu, " Fast and Parallel Algorithms For Solving Toeplitz Systems", Proc. ISMM, San Francisco, May 1981.

7. S.Y.Kung and Yu-Hen Hu, " A Highly Concurrent Algorithm and Pipelined Architecture For Solving Toeplitz Systems", Submitted to ASSP, 1982.

8. S.Y.Kung, " Matrix Data Flow Language for Matrix-Operation-Dedicated Array Processors", Proc. Europe Conf. Circuit Theory and Design, Hague, Netherlands, Aug. 1981.

9. S.Y.Kung, " Highly Parallel Architectures For Solving Linear Equations", ICASSP, Atlanta, 1981.

10. S.Y.Kung et.al.," A Matrix Data Flow Language/Architecture For Parallel Matrix Operations Based On Computational Wavefront Concept", Proc. CMU Conf. On VLSI Syst. and Computations, Oct. 1981.

11. S.Y.Kung et. al, " Wavefront Array Processor : Architecture, Language and Applications", MIT Conf. on Advanced Research in VLSI, Jan. 1982.

12. J.Goutsias and J.M.Mendel, " Parallel Processing Techniques For Generating Synthetic Siesmograms", SEO Report.

13. R.Chellappa and S.Y.Kung, "On Two Dimensional Markov Spectral Estimation", CISS, Princeton, Mar. 1982

NAVAL OCEAN SYSTEMS CENTER

14. H.J.Whitehouse and J.M.Speiser, " Sonar Applications Of Systolic Array Technology", Conf. IEEE Eascon, Nov. 1981.

DATE
ILMED
8